

A Channel Access Scheme for Large Dense Packet Radio Networks

Timothy J. Shepard*

BBN Systems and Technologies
10 Moulton Street, Cambridge, MA 02138

Abstract

Prior work in the field of packet radio networks has often assumed a simple success-if-exclusive model of successful reception. This simple model is insufficient to model interference in large dense packet radio networks accurately. In this paper we present a model that more closely approximates communication theory and the underlying physics of radio communication. Using this model we present a decentralized channel access scheme for scalable packet radio networks that is free of packet loss due to collisions and that at each hop requires no per-packet transmissions other than the single transmission used to convey the packet to the next-hop station. We also show that with a modest fraction of the radio spectrum, pessimistic assumptions about propagation resulting in maximum-possible self-interference, and an optimistic view of future signal processing capabilities that a self-organizing packet radio network may scale to millions of stations within a metro area with raw per-station rates in the hundreds of megabits per second.

1 Introduction

Multi-hop packet radio network technology offers an appealing prospect: that communication about a neighborhood or a metropolitan area might be a zero-cost commodity usable by all and capable of operation independent of any wired infrastructure. Indeed there are already companies offering spread-spectrum radios that operate in license-free parts of the spectrum that can provide point-to-point links over a few kilometers between buildings. These radios are operated without coordination and are purchased and installed by the users. Might this anarchy grow and become the predominant means of communicating about a metropolitan area? Can a large number (millions) of packet radio stations concentrated in a metropolitan-sized area operate and

provide any useful level of performance?

Designing such a large-scale packet radio system remains a significant engineering challenge covering many layers of the system, from the radio modulation and detection methods used at the lowest layers, to the network-layer routing and location of services in the higher layers. This paper reports results of an engineering design effort directed at the channel access scheme, which fits somewhere between the above-mentioned layers. The issues addressed include the growth of the noise level as the system scales, controlling access to the channel without any centralization of control, and strategies for routing. Scalability and decentralization of control are the primary concerns of this effort. (Mobility, however, is not a primary goal of this effort. The goal here is to design an alternative for running cables between buildings.)

The results reported here include: (1) a new framework for analyzing the performance of large-scale packet radio systems, (2) an analysis of the decline of signal-to-noise ratios as packet radio systems scale to millions or billions of nodes within a single metropolitan area, and (3) a channel access scheme (involving spread-spectrum and the timing of transmissions) that can ensure collision-free transfer of packets to nearby nodes while requiring only local coordination between the stations and only a single transmission to convey each packet.

This paper begins with a fundamental look at packet radio network modeling, and develops a new model by which we can understand under what conditions packets will be received successfully in a large packet radio system in terms of signal levels and the noise and interference levels. (We will treat all interference as equivalent to noise.) A model of noise growth as the system scales is then used to evaluate the potential for packet radio networks to scale to millions or billions of stations within a metropolitan area. With an understanding of what is required for successful reception and what noise levels can be expected, we then develop a design for station behavior including a collision-free channel access scheme. Simulations of small networks (consisting of only 100 or 1000 stations) were used to demonstrate the effectiveness of the channel access scheme and verify the analysis, and are reported briefly here (and more extensively in [18]). We conclude with a brief discussion of the potential performance of such a large-scale metropolitan-area packet radio network.

*This article includes work adapted from the author's thesis [18] that was submitted in July of 1995 at the Massachusetts Institute of Technology. At M.I.T., this work was supported in part by the Advanced Research Projects Agency under contract number DABT63-92-C-0002 and the author was supported in part by an AT&T-Bell Laboratories Ph.D. Scholarship sponsored by the AT&T Foundation.

2 Existing models of communication

An assumption commonly made since the earliest days of packet radio networks is that any overlap (in time) of multiple transmissions at an intended receiver prevents the successful reception of any packet transmissions involved in the overlap [1]. For some modulation and detection schemes, particularly those that were practical for digital radio transmission the early 1970s, this assumption is reasonable. From this assumption (that any amount of interference is fatal to packet reception) came much important work in channel access schemes, including a particularly notable application to wired (non-radio) local-area networks with Ethernet's CSMA/CD scheme.

The development of channel access schemes for packet radio networks using this simple model of interference coupled with simple models of propagation has continued into this decade. The most notable recent progress in this area is the MACA-MACAW-FAMA line of work begun by Karn in [9] and continued in [4], [7], and [6]. These schemes are based on a simple model of propagation and interference as all-or-nothing phenomena. In the spirit of the original ALOHA [1], they are asynchronous, and provide random access to the channel. With these techniques, if there is no other traffic and a packet arrives to be sent at some station, then that station can immediately begin the protocol for accessing the channel. If there is no contention, the transmission (and reception) of the packet and any associated control messages will then proceed without delay.

An important idea in multihop packet radio networks is that the channel can be spatially reused. Stations in different locations may make use of the channel simultaneously if they are sufficiently separated so that mutual interference will not prevent the transfer of the packets. Parallelism may offer potential performance gains if the interference from distant stations does not reduce the throughput achieved with each transmission by an amount greater than the amount of parallelism. Modeling the intricacies of radio interference is problematic at the system design level, and earlier work in multihop packet-radio performance analysis (e.g. [24], [11], and [23]) has always sought to simplify the criteria for which situations do and do not lead to successful reception of a packet.

Commonly the goal is to get to a nodes-and-edges view of the network where interference between stations not connected by an edge may be safely ignored. Then a scheme is devised where the network can ensure that for each reception that only one hearable station is transmitting. Then traditional (non-radio) packet-network design and performance analysis techniques may be brought to bear. A textbook method for ensuring non-interfering use of the channel is to assume system-wide synchronization and control, divide time into non-overlapping slots, and assign a compatible set of transmissions to occur in each time slot [3]. Work has progressed on methods of finding good assignments of transmissions to slots (e.g. [14]).

For large systems, this view and approach are problematic in two ways: (1) aggregate interference from distant stations is ignored (which might not be safe when there are many stations), and (2) a large system may be difficult to synchronize reliably (particularly if elements of it are to be capable of autonomous operation) and to reliably control (if there are system-wide dependencies on station geometry in the algorithms used to assign slots). We will see in Section 4 that aggregate interference will be considerable in a large packet radio system.

Spread-spectrum's anti-jam capabilities have long been recognized as a potentially valuable tool for handling interference in packet radio networks (e.g. [8], [13], and [19]). Spread-spectrum radio techniques allow a designer to treat, to some extent, interfering signals as if they were random (thermal-like) noise. To what extent this is true is discussed in [19] where it is suggested (on page 1109) that 5 dB of additional signal-to-noise margin more than required by the Shannon bound [16] is sufficient to achieve a reasonably low error rate. This appears to be consistent with the results in a more detailed treatment offered in [22]. Treating interference as noise can not only greatly simplify the task of the designer of channel access schemes, but may also enhance the efficiency of a packet radio network by enabling greater parallelism through greater spatial-reuse [10]. Digital spread-spectrum radio techniques are just today becoming practicable due to advances in VLSI circuitry. In this work, we will attempt to take good advantage of spread-spectrum's abilities in order to distribute the control of the network as much as possible.

3 A new model

Here we develop a model for communication in a packet radio network, based on the theoretical Shannon bound, that better reflects the underlying realities affecting radio system performance. The important difference between this model and the prevalent all-or-nothing model is that the presence of a signal (including those from interferers) at a receiver is a phenomena that can occur with varying degree. First we will present a very general model of propagation, and then develop our model of reception in packet radio networks based on a simplification of the general model of propagation.

A model of communication between stations within a packet radio network can be viewed as having three parts: (1) a model of a signal (used to model received and transmitted signals); (2) a model of propagation; and (3) a model for determining when transmissions are successfully received by the intended recipient.

3.1 Signals

A signal (transmitted or received) is most completely modeled as a real-valued function of time. The signal transmitted by station i is denoted by $s_i(t)$. The received signal at station i is denoted by $y_i(t)$. At times when a station is not transmitting a message, its output signal is zero. When one station wishes to convey a message to another, it transmits a signal that influences the signal received by the receiver. The receiver attempts to recover the message from the received signal. The details of designing signals to carry messages and designing receivers to detect messages (from received signals) are beyond the scope of this paper.¹ Here we just need to understand what performance can be realistically achieved, and what parameters determine performance.

The two parameters of a transmitted signal that are important for understanding system performance are the transmitted signal's power level and its bandwidth. Both of these parameters are limited by government regulation and by the limitations of the particular transmitter hardware. Transmitter power will be assumed to be controllable. (Bandwidth of the radiated signal will most likely be fixed at time of manufacture.) Successful reception of the message at the receiver will depend upon receiver performance, the total

¹A good introduction and theoretical treatment may be found in [12]. For a thorough coverage of spread-spectrum techniques, see [19].

power level of interfering signals at the receiver, and the received power level of the signal containing the message.

3.2 A general model of propagation

The influence transmitted signals have on received signals is determined by *propagation*. Propagation and noise determine the received signals as a function of the transmitted signals (including interfering sources). Assuming linearity and time-invariance, a general model is

$$y_i(t) = n_i(t) + \sum_{j=1}^M h_{ij}(t) * s_j(t) \quad (1)$$

where M is the number of stations, $h_{ij}(t)$ is the response at station i to an impulse in time transmitted by station j (the $h_{ij}(t)$ will be collectively referred to as the propagation matrix H), $n_i(t)$ is the signal due to thermal noise at station i , and the symbol $*$ represents convolution. This model for propagation is not much of a simplification of the real world; the world between antennas is mostly linear and time-invariant. (We are not designing for mobility here, but even if we were, the variations due to station movement occur at a timescale vastly different than that over which t varies in the $h_{i,j}$ parameters.) The impulse response $h_{ij}(t)$ is a general model for propagation in that it can represent the strength of the propagation, the propagation delay, and any multi-path propagation.

3.3 Simplification of propagation model

We now simplify the model by ignoring both propagation delay and multi-path effects. Propagation delay and multi-path propagation are important effects to consider when actually designing the system, but their effects can be safely ignored here. If necessary, actual delays could be observed and easily compensated for in the scheduling technique presented in Section 7. The successful detection of wide-band spread-spectrum signals (which we will use) is particularly robust to interference from multi-path propagation. If necessary, a *rake receiver* [19] can be employed to detect and combine the separately arriving copies of a transmitted signal. We can expect that the reduction in performance due to actual multipath would be equivalent to a couple of decibel decrease in signal to interference ratio, which would only affect our performance conclusions by a small constant factor.

With this simplification, the $h_{ij}(t)$ are now just scalar multiples of the unit impulse, $h_{ij} \cdot \delta(t)$, and Eq. 1 can be simplified to

$$y_i(t) = n_i(t) + \sum_{j=1}^M h_{ij} s_j(t) \quad (2)$$

where the h_{ij} are now scalars instead of functions (so the received signals are modeled as noise plus a weighted sum of the transmitted signals).

3.4 Reception

Whether or not a given packet transmission will be successfully received in a real network will depend upon many technical details. However we can derive a bound on the performance of the receiver from Shannon's capacity theorem [16] if we assume that the receiver makes no attempt to model

and subtract the interfering signals. This assumption is reasonable given the number of interfering signals expected in a network of many stations. Techniques for multiuser detection that do estimate the interfering signals can surpass the bounds derived from Shannon's capacity theorem, but are practical only when there are few interfering signals.²

Shannon's theorem bounds the capacity C of a communication channel by a function of the average signal power S , the average interfering noise power N , and the bandwidth W :

$$C \leq W \log_2 \left(1 + \frac{S}{N} \right). \quad (3)$$

This bound can be used to provide a model for successful reception of a transmission if we assume that the transmitter and receiver are attempting to optimize the probability of successful reception. Achieving the Shannon bound is not practically possible, but for a reasonable effort, a capacity can be achieved that corresponds to the Shannon bound for a situation a few dB worse in signal-to-noise ratio.

In general, stations might vary the rate at which they communicate depending on the observed interference. This work will assume that all the stations will communicate at some rate that is fixed by the design, and will address what this rate should be.

A packet will be successfully received at a station i from station k if, while it is being received, the received signal-to-noise ratio is at least some small factor, $\gamma > 1$ (and probably around 3 which is equivalent to the 5 dB mentioned above), more than the minimum required signal-to-noise ratio, i.e.

$$\frac{S}{N} \geq \gamma \left(2^{\frac{C}{W}} - 1 \right). \quad (4)$$

C is now not exactly the capacity, but the data rate at which the stations are attempting to communicate. The signal strength S is the power of the signal received at station i from the sending station k (i.e. the power in the signal $h_{ik}s_k(t)$) and N is the power contained in the sum of the interfering signals,

$$N = \text{Power} \left(n_i(t) + \sum_{j=1, j \neq k}^M h_{ij}(t) s_j(t) \right). \quad (5)$$

The power in this signal is the same as the sum of the powers of each of the interfering signals, as we have assumed that the signals are uncorrelated and of zero mean. As we will see in Section 4, in a large system the interference from other stations will dominate any thermal noise, so the thermal noise may now be ignored. Hence the signal-to-noise ratio at a receiver i for the transmission from station k can be computed (for purposes of simulation) from just the powers of the transmitted signals and the h_{ij} 's squared,

$$\left(\frac{S}{N} \right)_{ik} = \frac{h_{ik}^2 P_k}{\left[\sum_{j=1}^M h_{ij}^2 P_j \right] - h_{ik}^2 P_k}. \quad (6)$$

The P_i 's, and hence S and N , are actually all functions of time and vary as stations begin and end transmissions and

²Verdú in [27] suggests that multiuser detection might be possible when the number of interfering signals does not exceed 10 to 15, and states that the complexity of multiuser detection is exponential in the number of interfering signals. The packet radio networks considered here might nevertheless benefit from receivers that model and subtract only a few of the strongest interfering signals, but consideration of this potential improvement of receiver performance will remain beyond the scope of this work.

vary transmitter power levels. The criterion for successful reception of a packet is then that the signal-to-noise ratio be greater than the required minimum for the duration of its reception, which can be determined from the power levels alone.

3.5 Calibration

Our propagation model is not complete until the h_{ij} are specified. In the real world, stations may observe the actual propagation between stations that are capable of direct communication. Precisely modeling propagation between each pair of stations not capable of direct communication will not be important. In this work, propagation will be modeled by setting each h_{ij} proportional to $1/r_{ij}$ where r_{ij} is the distance from station i to station j . This model corresponds to the familiar $1/r^2$ free space loss (in power) for electro-magnetic radiation.³

Actual propagation in most cases will either be nearly equal to the free space propagation (when the antennas are within radio line of sight) or will be attenuated (when there are obstructions). Hence, assuming free space propagation, we will tend to accurately model the strength (at a receiver) of the stronger signals from nearby sources while overestimating the strength of the many weaker signals from more distant transmitters.

4 Modeling noise growth as a system scales

Interference from both near and distant stations will affect the ability of a station to successfully receive packets. In later sections we will show how to manage interference from local sources explicitly. In this section, we will attempt to estimate the aggregate interference from many non-local stations. No precise distinction is made here between local and non-local. Nevertheless in later sections we will be able to roughly fit together our understanding of the level of interference from distant stations with our scheme for handling interference from local sources.

Spread-spectrum techniques of modulation and detection provide an ability to communicate in the presence of interference. But spread-spectrum methods are not cure-alls as the achievable data rates are bounded by the Shannon limit. In the presence of high levels of interference, signal-to-noise ratios will be reduced, and hence the rate of communication will be reduced. If the signal-to-noise ratios sink as the system scales, then the communication rates must sink as well.⁴ This section will examine how the signal-to-noise ratios decline as the system scales.

Pessimism must be a guide when attempting to evaluate the level of interference in a system of stations that will be deployed without the chance to engineer the number or the placement of stations. But in this case, too much pessimism leads to the conclusion that the whole scheme is unworkable. We will first examine a model with a bit too much pessimism, and then refine it slightly to return to a workable level of noise.

³The *power* of a signal captured by an antenna is proportional to the power per unit area of the incident electro-magnetic radiation. The power per unit area falls off as the inverse square of the distance from the source. The voltage measured on the feedline is proportional to the square root of the power on the feedline. See [15] Chapter 4.

⁴The relationship between signal-to-noise ratio and the Shannon bound on communication rate is essentially linear when the signal-to-noise ratio is significantly less than one. I.e. $\log_2(1+x)$ in Eq. 3 is $\frac{x}{\ln 2}$ (approximately 1.44x) when $x \ll 1$.

Assume that stations are distributed at some average density ρ throughout the infinite plane, and that each station is operating its transmitter at unit power output and at duty cycle η . Then the power radiated per unit area in the plane is (on average) $\eta\rho$. For a receiver located in the plane, the power level received Γ relative to the power received from a station at a distance of one characteristic length $R_0 = \rho^{-\frac{1}{2}}$, can be computed by integrating:

$$\Gamma = \int_0^\infty \frac{1}{r^2} \eta\rho 2\pi r dr. \quad (7)$$

Unfortunately, with the infinite upper bound, the integral diverges (regardless of what lower bound we might chose). Hence, for a receiver located in an infinite plane with a uniform and finite density of transmitters, the received power level would be infinite. The signal-to-noise ratio would be zero regardless of the source of the signal, so no communication would be possible.⁵ There are a number of ways out of this conundrum. The key is to notice that the integral just barely diverges. For example, the slightest bit of atmospheric attenuation, which would introduce an $e^{-\alpha r}$ factor to the integrand, would make the integral converge to a constant. Nor does the integral diverge if we integrate out to some reasonable bound, stopping short of infinity.

Fortunately, we do not live on an infinite flat earth. In UHF and higher bands, only stations that are not hidden over the horizon can contribute to the interference at a receiver.⁶ Hence the population of stations that are able to interfere with a given receiver will be limited to those in the same geographic region. If the earth's surface were perfectly spherical and all antennas were at the same height, then this region would be the interior of a circle. This limit on propagation *is* well modeled by a "transmission radius", but it is not a radius that can be engineered into the station or controlled as the system operates, but rather a consequence of station placement and surrounding terrain. A metropolitan area on flat terrain (or nestled in a bowl-shaped valley) may have all stations within direct line-of-sight propagation, hence the circle could cover at least an entire metropolitan area. The model for propagation is then $\frac{1}{r^2}$ within a circle encompassing a metropolitan area, with no interference from any stations outside the circle.

The growth in the overall level of interference as the system grows in number (and density) can now be estimated. Assume that M interfering stations are distributed randomly within a circle of radius R , and that stations outside the circle can be ignored. The average density ρ is then $\frac{M}{\pi R^2}$. As the number of stations M increases, so does the density. The distance to nearest neighbors also decreases, remaining proportional to the distance $R_0 = \rho^{-\frac{1}{2}}$. The signal level S from such a nearest neighbor transmitting with

⁵This observation is similar to a troublesome answer to the question "Why is the sky dark at night?" (Olbers' paradox). If we assume an infinitely large and infinitely old universe, a constant density of galaxies in the universe, and that each galaxy radiates a given amount of power, and then perform a similar integration (this time in three dimensions) we can conclude that an infinite amount of power should be impinging upon our eye when we look up at the sky. Presumably we have made an assumption that is not true for the universe in which we reside.

⁶The term *horizon* here means something slightly different at radio frequencies than it does at light frequencies. At radio frequencies there is a significant amount of propagation over the visual horizon, but a horizon-like attenuation of signals occurs nevertheless at a distance further out. The distance at which this attenuation occurs is sometimes called the *radio horizon* and is often modeled as if it behaved like a visual horizon of an earth with the radius increased to $\frac{4}{3}$ of the actual earth's radius. For more see [15] Chapters 3 and 4.




Figure 1: Decline of the signal-to-noise ratios as M , the number of stations, grows (Eq. 15). Each member of the family of curves is for a different value of the duty cycle, η (denoted as “eta” in the curve labels).

unit power would be

$$S = \frac{\alpha}{R_0^2} \quad (8)$$

$$= \frac{\alpha}{\left(\frac{1}{\sqrt{\rho}}\right)^2} \quad (9)$$

$$= \alpha\rho \quad (10)$$

where α depends on the antennas and wavelength used. The total power of interfering signals, N , ignoring the contribution from local interference inside the circle⁷ of radius $R_0 = \rho^{-\frac{1}{2}}$ (and ignoring other sources of noise) can be calculated as

$$N = \int_{R_0}^R \alpha \frac{1}{r^2} \eta \rho 2\pi r dr \quad (11)$$

$$= \alpha \eta \rho 2\pi \ln \frac{R}{R_0} \quad (12)$$

$$= \alpha \eta \rho \pi \ln \frac{M}{\pi}. \quad (13)$$

So the signal-to-noise ratio (SNR) is

$$\frac{S}{N} = \frac{\alpha\rho}{\alpha \eta \rho \pi \ln \frac{M}{\pi}} \quad (14)$$

$$= \frac{1}{\eta \pi \ln \frac{M}{\pi}}. \quad (15)$$

Thus the expected signal-to-noise ratio of a signal from one of the nearest neighbors depends only on $\ln M$ (the log of the total number of stations) and η (the duty cycle), and is independent of scale-length. Figure 1 shows a plot of the log of the signal-to-noise ratio as a function of the base ten log of the number of stations. The signal-to-noise ratio falls very slowly, approaching -20 dB for $\eta = 1$ as the number

⁷We cannot just simply drop the lower bound of this integral to zero for then the integral would blow up. But interference from local sources will be managed separately and explicitly later. Choosing a lower bound of $\rho^{-\frac{1}{2}}$ is reasonable (as stations closer than this distance are clearly local) and convenient (because it makes the algebra work out nicely).

of stations approaches 10^{12} . This observation is encouraging. The signal-to-noise ratio of a neighbor’s transmission falls slowly even as the number of stations grows exponentially (even with $\eta = 1$, it does not reach -21 dB until 10^{18} stations).

According to this model, for almost any realistically large population of stations in a packet radio network, direct communication at a definite rate with nearby neighbors (neighbors nearer than $\rho^{-\frac{1}{2}}$) should remain possible, provided that the stations can cope with signal-to-noise ratios of around -20 dB. Indeed, by Shannon’s capacity theorem, $C = W \log \left(1 + \frac{S}{N}\right)$, even with a signal-to-noise ratio of one part in one hundred, the theoretical communication capacity remains non-zero. In this case, $C = W \log_2(1.01)$, thus $\frac{C}{W} = 0.014$, or theoretical capacity of approximately 14 bits per second per kilohertz of channel bandwidth.

Thus far, these calculations have assumed $\eta = 1$. Stations will have to spend at least some of the time listening. For more reasonable values of η , the noise levels are improved. At an average duty cycle of one quarter, $\eta = 0.25$, the signal-to-noise ratio is better by a factor of four, or $+6$ dB. The resulting signal-to-noise ratio of around -14 dB yields a theoretical capacity of around 56 bits per second per kilohertz of channel bandwidth, but only when the station is transmitting. There is no gain in throughput by further reducing the transmit duty cycle in a large noisy system. Halving the duty cycle increases the average signal-to-noise ratio by a factor of two, which improves the data rate (while transmitting) by approximately a factor of two, but would result in no net gain in performance since the transmitters would then be operating for only half of the original amount of time.

What about neighbors that are not so near? The placement of stations or the routing algorithm might require direct communication between stations that are more than $\rho^{-\frac{1}{2}}$ distance apart. Free-space radio propagation falls off by a factor of four, or -6 dB, for each doubling in distance, so we can expect that a station at a distance of $2\rho^{-\frac{1}{2}}$ will be heard with a signal-to-noise ratio reduced by a factor of four, or -6 dB. Another factor of two in distance would be another -6 dB. Each 6 dB reduction in signal-to-noise ratio reduces achievable throughput by a factor of four. Thus, in large scale packet radio networks, direct communication (at a reasonable rate) will be possible only with nearby neighbors.

While this model gives us an estimate of the overall noise levels, the exact value of the signal-to-noise ratio will depend on the details of station placement and transmission control. There remains the problem of managing the transmissions on an individual and local basis. For example, interference from a very near station might amount to much stronger interference than the aggregate interference produced by distant stations, but since a nearby station is local, the interference can be managed locally.

5 Collisions

In more conventional models of packet radio networks (those involving a hard transmission radius and simple success-if-exclusive criterion for successful reception) the term *collision* is often used to describe how packets are lost due to interference. In the more sophisticated model used in this work (where the criterion for success is sufficiency of the signal-to-noise ratio at the receiver) the term collision may be misleading in that it suggests an overly simple interpre-




Figure 2: Examples of each type of collision. The X indicates the lost packet.

tation of the interaction between packets at receivers. In the model in this work, whether or not a packet is received successfully depends on more than just the number of simultaneous signals at a receiver. Nevertheless, a taxonomy of collision types will help us to understand local interference, even in the context of our more complete model.

If a collision occurs, then it must fall into one of the following three cases (see Figure 2):

1. Collisions due to the transmission of another packet from a station not involved in the exchange of the dropped packet, which is not addressed to the station receiving the dropped packet.
2. Collisions due to multiple stations attempting to send packets simultaneously to a single station.
3. Collisions due to a packet arriving at a station while another packet is being transmitted by the receiving station.

This enumeration covers all possible cases of an interfering transmission. If the interfering transmission does not involve the receiving station, either as a receiver or transmitter, then it is a Type 1 collision. If it does involve the receiving station as the intended target of the interfering transmission, then it is a Type 2 collision. If it involves the receiving station as the sender of the interfering transmission, then it is a Type 3 collision. Multiple collision types may occur simultaneously in more complicated situations.

Our use of spread spectrum can eliminate most packet loss due to Type 1 collisions. If a nearby interfering station is transmitting, and the receiver is already prepared to cope with a signal-to-noise ratio of $\frac{1}{100}$ due to the (potential) overall level of noise, then in order to significantly increase the level of interference, a nearby station would have to be very near indeed. Even a station at one fourth of the $\rho^{-\frac{1}{2}}$ distance would have only a small effect on the total amount of

interference.⁸ So, in most cases, this type of collision is not a problem. When stations are so close that Type 1 interference is a problem, then it really is a local problem, and the stations involved must cooperate and each must refrain from transmitting in a manner that interferes excessively with the receptions at its neighbor. A method of achieving this coordination in a decentralized manner will be presented in Section 7.

Type 2 collisions are very similar to Type 1 collisions. The only difference is that the intended receiver of both transmissions is the same station. These can be eliminated by enabling stations to receive multiple transmissions in parallel. With spread-spectrum radio receivers, elimination of packet loss due to this type of collision requires only multiple tracking and despreading channels. A multiplicity of despreading channels is already a common feature of existing spread-spectrum receivers. For example, GPS (Global Positioning System) receivers often have six or twelve despreading channels. With a sufficient number of despreading channels, packet loss due to Type 2 collisions can be eliminated. The number of despreading channels needed will depend on the details of the routing and transmission control schemes used, but in any case, it should not be larger than the number of neighbors that might communicate directly with the station. This number should be small, since, as we have already seen, only nearby stations will be capable of direct communication over the din of background noise. (A routing strategy that will be presented in the next section was used in a number of simulations of randomly placed stations and the number of routing neighbors never exceeded eight.)

Type 3 collisions are a more difficult problem. The interference from a transmitter located with a receiver will be so powerful that no feasible amount of processing gain (even when combined with the isolation provided by the antenna duplexer) can achieve reception while the local transmitter is operating. But like the nearby case of Type 1 collisions, packet loss due to Type 3 collision is a local problem, and need not be solved globally. It is sufficient to ensure that the local transmitter does not operate at times when other stations might send a packet to the local receiver. A method of achieving this coordination in a decentralized manner will be presented in Section 7.

6 Design strategy

A design strategy for a viable large-scale packet radio network can now be described. The analysis in Section 4 showed that if there are many (millions of) stations in an area, the stations will be immersed in a din of interference. Nevertheless, by using spread-spectrum radio techniques with a moderately high processing gain (in the range of 20 dB to 25 dB) stations will be able to communicate directly with nearby neighbors (those stations within a distance of approximately $2\rho^{-\frac{1}{2}}$) even as the system scales to a large number of stations. By using spread spectrum, the interference from distant stations can be treated as random noise, and no system-wide coordination is needed to manage the use of the channel.

By cooperatively forwarding packets, the stations may organize themselves into a fully connected network to allow communication beyond the immediate neighbors. Whether

⁸ A low-power signal added to a high-power signal yields a signal with power level not much different than that of the original high-power signal.

or not the network is fully connected will depend on if there are enough stations to blanket the area, if the stations are distributed uniformly enough, and what distance can be covered in one hop. The analysis that produced Figure 1 assumed that the neighbor was $\rho^{-\frac{1}{2}}$ distance away and that the interfering transmissions were at the same power level and were evenly distributed throughout the region. But the design will need to accommodate communication between neighbors that are not exactly at this distance, and will need to cope with varying densities.

If the stations are distributed randomly and independently in the plane at density ρ , and if $\rho^{-\frac{1}{2}}$ is the maximum distance that can be covered in a single hop, then the expected number of neighbors that a station will have is the expected number of stations inside a circle of radius $\rho^{-\frac{1}{2}}$, which is $\rho\pi\left(\rho^{-\frac{1}{2}}\right)^2 = \pi$. That the expected number is only π suggests that $\rho^{-\frac{1}{2}}$ may not be a far enough reach to ensure connectivity. No claim about connectivity can be made without knowledge of the particular geometry, but it is reasonable to expect that variations in density will at some stations require reaching farther than to just three expected stations. Doubling the distance to $2\rho^{-\frac{1}{2}}$ (by increasing the processing gain by 6 dB or a factor of four) should suffice in most situations. Assuming again uniform distribution, the expected number of reachable stations would then be 4π . This doubling of range comes at the expense of a factor of four in raw throughput since the processing gain has been increased, and any further increase in range (by increased tolerance to interference) would further impact throughput. (A doubling in range would quadruple the noise-to-signal ratios, reducing raw throughput by a factor of four since achievable throughput depends linearly on signal-to-noise ratio in a noisy system.) With the signal-to-noise ratios for stations at $\rho^{-\frac{1}{2}}$ distance in the -10 dB to -15 dB range for reasonable duty cycles, the need to budget around 5 dB of headroom for successful detection in the receiver, and the need for an additional 6 dB margin for more distant stations, the proper amount of processing gain is determined to lie in the range of 20 to 25 dB.

6.1 Power control

In the analysis so far, all transmissions were assumed to be at the same power level. In cases where stations are closer than maximum range, transmitting at full power is excessive. A more sensible approach is to control the power. If the stations are controlling power but are still transmitting with the same average power density as before, then the analysis of average signal-to-noise ratio remains the same. But by reducing power in situations where lower power levels can still deliver a sufficient signal-to-noise ratio at the intended receiver, interference to other stations can be reduced, increasing the signal-to-noise ratios in receivers at other stations.

The ratio of the average noise power level to the average signal level should not be affected by power control. This criterion strongly suggests a power control algorithm: *transmit with sufficient power to deliver a constant pre-determined amount of power to the intended receiver.*⁹ The choice of

⁹A better idea might be to transmit with power sufficient to just achieve the necessary signal-to-noise ratio. That would require knowing what the noise levels at the receiver will be, but the recent past might be a good-enough predictor of the future noise levels. This idea will not be explored further here.

pre-determined power level is not critical, because increasing or decreasing it will just slide all power levels in the system up or down, maintaining the same ratios everywhere, including the received signal-to-noise ratios. By fixing the received power level, the variance in signal-to-noise ratio can be reduced.

This power control algorithm is also appealing for another reason: as different areas in a network may vary in density, stations will automatically compensate by controlling power levels to deliver the same amount of power to the intended receiver. The average power density then remains roughly constant: if the density in some area is quadrupled, the distance to neighbors is cut in half, so power levels can be cut by a quarter, maintaining constant power density as the station density varies. Therefore the analysis from Section 4 remains applicable even to networks employing this method of power control.

6.2 Minimum-energy routes

We already know that packets traveling more than $2\rho^{-\frac{1}{2}}$ must be routed through intermediate stations. When there is a candidate intermediate station, and an option exists of either sending the packet directly or through the intermediate station, which should be done? In some sense, with power control, this choice will always exist, for if we choose to send the packet directly, we can increase the transmitted power to deliver the proper amount of power to the intended receiver. But if stations routinely did this to communicate directly with distant stations (stations significantly farther than $2\rho^{-\frac{1}{2}}$ where ρ is in this case the density in the immediate area), then we would be violating a crucial assumption of the earlier analysis: that the power density is constant and roughly $\rho\eta$. Violating this assumption in this way would significantly reduce the signal-to-noise ratios. Such high-power transmissions would also cause a high level of interference to the (presumably numerous) neighbors close to the transmitter. The criteria used to determine routes will need to prefer the short hops, which produce less interference, and avoid skipping over intermediate stations.

In an actual network, the stations may not know where they are geometrically, but they will be able to observe the path gains between themselves and construct entries in the propagation matrix H for the hops that are usable. A criterion for selecting routes that is directly determinable from the propagation matrix would be particularly convenient. A routing criterion that is directly determinable from the propagation matrix and that seems to meet our needs is *minimum-energy routing*.¹⁰ Consider the scenario in Figure 3. Station A wishes to send a packet to station C. Station B is a candidate intermediate station. Using minimum-energy routing, station B should be used as an intermediate hop if it reduces the total amount of interference to a distant station D *caused by the movement of this packet*. If B is used as an intermediate hop, the duration of the interference (to station D) caused by this packet will be doubled, but the power levels of the two transmissions may be much less than the single hop transmission. They would be less by as much as a factor of four if station B is exactly centered between stations A and C. Taking this intermediate hop would reduce the level of interference by as much as a factor of four at any distant station (e.g. D). Then the total energy (power integrated by time) of the interference to a

¹⁰The idea of minimizing energy is mentioned in [10] and credited there to David Mills.



Figure 3: Situation where there is a candidate station B for an intermediate hop between A and C. Routes are chosen so as to minimize each packet's total contribution to interference at distant stations (e.g. D).

distant station caused by this packet will be reduced by a factor of two.

Geometrically, with $\frac{1}{r^\alpha}$ free-space loss, minimum-energy routing will always take the intermediate hop if it lies within the circle which has a diameter with endpoints at Station A and Station C (i.e. the smallest circle that just touches both Stations A and C). Thus minimum-energy routing will choose routes that respect the local density and will not skip over intermediate hops. By using minimum-energy routing, the interference from the ensemble effect of many packets traversing the network is kept to a minimum, enabling as much raw data throughput as possible across each local hop. There are trade-offs. For example, this approach does not minimize latency. The multitude of store-and-forward delays incurred by always taking intermediate hops will adversely affect delay. This routing method may be inappropriate if delay is the overriding concern.

Minimum-energy routes are straightforward to compute, at least in networks of moderate size. The common algorithms for computing min-cost paths in networks¹¹ can be used to find the least-cost paths in the propagation matrix H , where the costs are the reciprocal of the path gains. (The reciprocal of the path gain is proportional to the power that would be used with power control.) The algorithm is also easy to distribute. Each station need only remember the next hop for each potential destination and the total energy along that route to the destination. Hop-by-hop routing is possible since, at each station, each transit packet will be routed as if it had originated at the transit station. In other words, a minimum-energy route from A to C that goes through B will use the same route from B to C as any other route that goes through B to get to C.

In large networks, routing remains challenging. We have identified here one criterion by which effective routes may be chosen (minimization of energies). We have not identified a method of computing these routes that can of scale to a network with millions or billions of haphazardly-placed stations. This problem, along with other higher-level issues (such as location of services), remain topics for further research. One approach to this challenge is described in [26].

7 Collision-free channel access scheme

When a packet is to be sent to another station, it must be sent at a time when its reception will not be prevented by the level of interference. If (as discussed in Section 5) we are only concerned about interference from the receiving station's own transmitter, then meeting this constraint can be easily accomplished: the sending station only needs to know

¹¹For example, the Distributed Asynchronous Bellman-Ford Algorithm is described in [3].

at what times the receiving station may be transmitting. If the receiver's schedule is known by the sending station, then the sending station can choose to send the packet at some other time. In order to make its schedule known to its neighbors, a station (the intended receiver) needs to schedule the times that it may be transmitting and inform its neighbors what those times will be.

Global clock synchronization is not required. Only the ability to relate one station's clock¹² with another's is required. This ability can be accomplished if stations occasionally rendezvous and exchange clock readings. Differences between clocks and small differences in clock rates can be mutually modeled, and the resulting models, along with the published transmit schedules (each reckoned by the publishing station's clock), can be used by neighbors to predict when a station will be transmitting.¹³

The method is as follows. Each station will independently produce and publish a schedule for itself. A schedule divides time into receive windows and transmit windows for a station. The schedule published by a station is a commitment by that station to listen (refrain from transmitting) at particular times in the future (during the receive windows). A station with a packet to be sent to another station will compare its own schedule with the receiving station's schedule and send the packet during a time when one of its own transmit windows overlaps with a receive window of the receiving station enough to handle the packet length. Each station only needs to be aware of the schedules of the immediate neighbors to which it might be directly sending packets.

The schedules must be devised so that the overlaps will exist for all pairs of stations that may communicate directly. Simple periodic schedules will not do. If two stations using simple periodic schedules were to happen to be running at the same phase, then communication between them would not be possible. (Choosing clock values to avoid unfortunate phase offsets is not possible if the stations are to produce the schedules independently without any global coordination.) This problem is solved by using schedules that are random or pseudo-random. If each station independently chooses a random schedule, then these schedules will allow many opportunities to communicate between any pair of stations.

7.1 Unaligned slots

One possible method of implementing the pseudo-random schedules is presented here. Implementing the scheduling method requires a method of generating the random schedules and a method of communicating the schedules and the clock readings to neighbors. If each station's clock is set differently, then the stations can all use a single schedule, each reckoned by its own clock. With all stations using the same schedule, then only the clock readings need to be communicated between stations. Time can then be divided into equal size slots, again reckoned independently by each station's clock, and each slot designated to be either a transmit or receive slot.

The schedule (singular since all may be the same) is a binary-valued function of a clock reading that divides time into transmit windows and receive windows. The schedule

¹²The term *clock* as used in this work does not imply knowledge of what time it is. Here clock just means something that advances at some known rate.

¹³See [25] for an example of how the drift of a clock driven by a quartz oscillator can be modeled from historical data and for a demonstration of how the model can then be used to accurately predict future drift.




Figure 4: A sample pseudo-random schedule for 20 stations. The line is drawn in slots where the station is scheduled to transmit, and the line is omitted in slots where the station is scheduled to listen. The receive duty cycle (the average fraction of slots that are scheduled for listening) here is 0.3. To send a packet from one station to another, it must be scheduled to fit in a period of time when the sending station is allowed to transmit and when the receiving station is listening. For example, at the circled time above station zero could not send to station 1 or station 2, but could transmit to station 3.

will be divided into equal-length slots of length T_{slot} , with all times in a slot sharing the same value (transmitting or receiving). Whether a particular slot is for transmitting or receiving can be determined by using a hash function to hash the value of time at the beginning of the slot. If the hash value is less than a threshold, then the slot is a receive slot, otherwise, it is a transmit slot. The threshold is selected to achieve the desired duty cycle. The receive duty cycle, p , is the probability that a slot is a receive slot.

The slots are unaligned, as is shown in Figure 4. Between stations there is no synchronization of clocks and the slot boundaries at a station are determined by the station's own clock. A slot at a station will overlap with parts of two slots of a neighbor, unless that neighbor's clock happens to be exactly aligned. The amount of overlap (or phase difference) between two stations is random and remains constant except for any drift caused by differences in clock rates.

Each station needs to set its clock in a way that ensures that it is set differently than each of its neighbors with which it will be directly exchanging packets. If the clocks were not set differently, then the identical schedules would prevent communication between the two stations. Clocks with only a small difference (of less than one slot time) would not have the full expected amount of time available between them to communicate as their transmit schedules would be somewhat correlated. But if there is at least one slot's time difference between the two clocks, then the schedules will be uncorrelated and we can treat each station's schedule as random and independent.

A simple way to set the clocks so that they are different is to set them independently to a random value. The probability that a station's clock may by chance be set to a value that is close enough to the value of neighbor's clock to cause trouble can be made arbitrarily small by increasing the number of significant high-order bits in the clock. Each additional high-order bit added and initialized randomly will reduce the probability of such an unfortunate coincidence by a factor of two.

7.2 Performance of scheduling scheme

As was discussed at the end of Section 4 the throughput performance of a large packet radio system is to first order independent of the duty cycle at which all stations are operating. In [18] the parameters of this scheduling method are explored and a 30% receive-duty cycle is found to be nearly-optimal for a wide range of situations. So an individual station may spend about 70% of the time transmitting, and the expected fraction of time at which transmission is possible to any individual neighbor is 21% (30% of 70%). In [18] is also presented a method of scheduling the packets into the slots by limiting the packets to a small fixed-size one-fourth the length of a slot time. This further limits the amount of time to an individual neighbor to 75% of the total time when transmission is possible, or approximately 15% of all time. Note that a station may achieve close to 15% of its raw rate with two (or a few) different neighbors simultaneously since little of the time available to send to one neighbor will conflict with the time available to send to another neighbor.

The additional delays due to the scheduling scheme are fairly well modeled by a Bernoulli process with the Bernoulli trial probability of success of $p(1-p)$ which is equal to 0.21 for $p = 0.3$. Hence the expected number of slots until the packet can be sent is $\frac{1}{p(1-p)}$, which for $p = 0.3$ is 4.76 slot times. This assumes that no other traffic is contending for the stations transmitter. Even with other traffic, a station need not block the head of the line. Traffic to other stations may be transmitted while waiting for a suitable time to arrive. With no head-of-line blocking, stations may achieve transmit duty cycles approaching 50% (as demonstrated in [18]).

7.3 Respecting neighbors' receive windows

So far, we have developed a method of avoiding interference from a receiving station's own transmitter. Now we will extend this method to avoid all significant interference from local sources.

Interference from a nearby station's transmitter may be a problem if it is used to transmit at high power (to deliver a packet to a distant station). It would be a problem if the nearby station's transmitter delivers an interfering signal with power sufficient to significantly lower the signal-to-noise ratio of packet receptions. Whether the effect is significant or not will depend on how much processing gain the stations are using (which determines their tolerance to interference).

The power levels of signals are commonly discussed in terms of decibels, a logarithm of the power level. But here we are concerned with what the effect of an additional interfering signal (with some power level) is on the overall level of interference, which may already be quite high. The power levels add, but not the logarithms of the power levels. For example, if two signals, one at a power level of 20 dB and the other at a power level of 10 dB are added, the power level of the resulting signal is at 20.4 dB, which is a barely significant change. In order for the addition of a weak signal to increase the overall level of interference by more than 1 dB its power level must be at least one fourth the power level of the overall interference. One decibel, which is about a 25 percent change, is a reasonable threshold for significance. While we can not strictly budget the additional level of interference we may tolerate from each nearby neighbor independently (as two additional sources of interference can



Figure 5: Region of *respect neighbor* constraint on packet scheduling. A station would have to be inside the circle for a transmission from A to B to significantly raise the level of noise at its receiver. In those rare cases, Station A must respect (by not overlapping) the receive windows of any stations inside the circle when scheduling transmissions to Station B.

combine to produce an even greater level of extra interference), we can hold each such potential additional source of interference to a maximum increase of 1 dB in total interference and budget a few decibels of additional headroom. It would then take more than four simultaneous high-power transmissions (each contributing just under the 1 dB threshold) from nearby neighbors to have more than a 3 dB effect on the overall level of interference.

Only in infrequent circumstances will a neighbor's transmission increase the level of interference by more than 1 dB. In Section 4 it was shown that the background level of noise may be roughly a factor of 100 greater (up by 20 dB) than the level of individual signals received from nearby stations. Stations already must cope with this level of interference. In order for an interfering station to significantly increase (by more than 1 dB) the total amount of interference, it would have to deliver (to the interfere-ee) more than 20 times (or 13 dB more) the amount of power that it is delivering to the intended recipient. (If the noise level is 20 dB over the target receive power, then the threshold of significance of one fourth, or -6 dB, of the level of the noise is 14 dB over the target receive power. Choosing 13 dB here is one decibel more conservative.) Assuming $\frac{1}{r^2}$ propagation, this threshold will be exceeded only when the receiving station is more than five times as far away as the interfere-ee (from the transmitter). For example, in Figure 5 a station would need to be located inside the circle for a transmission from Station A to Station B to significantly affect its noise level. If the most distant stations we are communicating with are at a distance of $2\rho^{-\frac{1}{2}}$, then the expected number of stations inside a circle with a radius of one-fifth this distance is only $\pi \left(\frac{2}{5}\right)^2 \approx 0.5$. This number is well under the interference threshold of 4 nearby transmitters selected above. Therefore this form of interference will not often be a problem.

When high power must be used, an additional constraint can be placed on the scheduling (to avoid interfering with a neighbor's reception). Those packet transmissions that will require high power must not be scheduled at a time that overlaps with a receive window at a neighbor who is too close. Using the notation from Section 3, station k is too close to station i for station i to send to station j during a receive window of station k if $h_{i,j}^2 < 20h_{i,k}^2$. The factor of 20 means that if the signal at a nearby neighbor will, at a neighbor, be more than 13 dB above the usual reception power level, then it cannot be sent during that neighbor's receive windows. While in simulations [18] this constraint seldom came into play it succeeds in effectively eliminating




Figure 6: A network of 1,000 randomly-placed stations. Line segments connecting the stations show the direct hops used by minimum-energy routes.

packet receptions with low signal-to-noise ratios.

8 Simulation to observer SNR

The decentralized channel access scheme presented here has been demonstrated in simulation [18]. One-thousand stations were placed randomly in a square area, and minimum-energy routes were computed between each pair of stations. All the direct hops used by any route is shown in Figure 6. Figure 7 shows a distribution (over all packets at all stations) of the signal-to-noise ratios when this network is maximally-loaded (unlimited traffic available at every station waiting to be sent to each of its neighbors). The stations were on average transmitting about one third of the time (around a 30% duty cycle). This distribution is centered around -9 dB which agrees with Figure 1. That the signal to noise ratio never dropped below -17 dB during the reception of any packets demonstrates that all packets may be received successfully (by our model) if there is sufficient spread-spectrum processing gain (which in this case would need to be around 23 dB). This would allow around one bit per second of raw data rate for every 200 Hz of bandwidth. A more detailed discussion of the performance of systems using this scheme is presented in [18].

9 Conclusion

A single-channel packet radio system can scale to millions or billions of stations within a metropolitan area while maintaining the ability for each station to communicate with its nearby neighbors at a rate that will not continue sink appreciably as the system continues to scale beyond a million stations. With completely decentralized control, and only a single transmission to convey each packet, packets can be transferred to nearby neighbors without any loss due to collisions.




Figure 7: Distribution of worst signal-to-noise ratio during reception of each packet in a 1000 station network. Bottom plot is same as top plot but at a magnified vertical scale so that the left tail can be seen more clearly. No packets were received with a signal-to-noise ratio of less than -17 dB.

10 Final Thoughts

One promising improvement would be to take advantage of directional gain. For each decibel of directional gain achieved at both the transmitting and receiving station, two fewer decibels of processing gain would be needed. So a set of six directional antennas (each with a modest 6 dB of gain in a main lobe roughly $\pi/3$ radians wide) each with its own independent transmitters and receivers at each station would increase data rates in the system by more than a factor of ten (with 12 dB less processing gain required). So raw data rates of around 0.1 bits per second per Hertz of bandwidth would then be achievable to nearby neighbors in a large system.

In the U.S. the FCC is considering allocation of the 59 to 64 GHz millimeter-wave band to unlicensed and unregulated communication system (FCC docket number 94-124). There is an absorption band here due to resonance with a quantum transition of the O_2 molecule. It attenuates signals by as much as 15 dB/km over the usual free-space loss [5]. This band would be ideal for this sort of packet radio system. The attenuation would improve the signal-to-noise ratios in a large system by reducing the impact distant stations may have on the noise level at each receiver, though it would also reduce the maximum distance over which stations would be able to communicate in low-density areas. We can imagine that in less than ten years it will be feasible to build spread-spectrum systems in this band with spreading-code chip rates of around 2 GHz. In this band, and with 6 dB of directional gain, a packet radio system could have raw transmitter data rates of around 200 megabits per second, have transmitter duty cycles of 25% to 50%, and be scalable to millions of stations within a metropolitan area. (To visualize this, look at Figure 6 and think of each line connecting stations as representing a 25 megabit-per-second bi-directional link.)

The performance of such a system could rival that of traditional metropolitan area telephone systems. With a packet size of around 500 bits (or $2.5 \mu s$ at 200 Mb/s), slot sizes four times larger than a packet ($10 \mu s$), and expected scheduling delays at each hop of around 5 slot-times, the per-hop delay would be $50 \mu s$. The number of hops in a multihop path across a metropolitan area will be roughly the square root of the number of stations, or roughly a thousand hops in a network of one-million stations. Hence one-way delays to cross a metro area could be around 50 ms. Aggregate capacity would be fairly high. If the area were divided in half (right down the middle) then the number of ways across the divide would be roughly the square root of the number of stations. In a network of one million stations, that would be 1000 such crossings, and at 25 Mb/s each, that would be 25 Gb/s aggregate capacity from each half of the network to the other half (with of course additional capacity remaining within each half). 25 Gb/s divided by 500,000 stations is 50 kb/s per station.

11 Acknowledgments

I would like to thank all those who have helped me with the thesis and this paper, especially my thesis supervisor David D. Clark, committee members Thomas F. Knight Jr. and Jerome H. Saltzer, and many others at M.I.T. for their helpful comments while doing the research and writing the thesis. Craig Partridge helped motivate me to get started with the process of distilling this paper from my thesis, helped me outline the first three sections, and read a draft of the paper and provided useful feedback. Isidro Castiñeyra read a draft and provided some particularly useful feedback. I would also like to thank the anonymous reviewers for their helpful comments.

References

- [1] Norman Abramson. THE ALOHA SYSTEM—Another alternative for computer communications. In *Proceedings of the Fall Joint Computer Conference*, pages 281–285. AFIPS, 1970.
- [2] Norman Abramson, editor. *Multiple Access Communications*. IEEE Press, 1993.
- [3] Dimitri Bertsekas and Robert Gallager. *Data Networks*. Prentice-Hall, 1987.
- [4] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. MACAW: A media access protocol for wireless lan's. In *Proceedings of SIGCOMM'94*, pages 212–225. ACM, 1994.
- [5] CCIR Rep. 719-2. Attenuation by atmospheric gases. In *Recommendations and Reports of the CCIR*, volume V, pages 167–177. Geneva, 1986. CCIR, International Telecommunication Union.
- [6] Chane L. Fullmer and J.J. Garcia-Luna-Aceves. FAMA-PJ: A channel access protocol for wireless LANs. In *MobiCom'95 Proceedings of the First Annual International Conference on Mobile Computing and Networking*, pages 76–85. ACM, November 1995.
- [7] Chane L. Fullmer and J.J. Garcia-Luna-Aceves. Floor acquisition multiple access (FAMA) for packet-radio networks. In *Proceedings of ACM SIGCOMM'95*, pages 262–273, 1995.

- [8] Robert E. Kahn, Steven A. Gronemeyer, Jerry Burchfiel, and Ronald C. Kunzelman. Advances in packet radio technology. *Proceedings of the IEEE*, 66(11):1468–1496, November 1978.
- [9] Phil Karn. MACA - A new channel access method for packet radio. In *ARRL 9th Computer Networking Conference*. ARRL, 1990.
- [10] Phil Karn. Spectral efficiency considerations for packet radio. In *ARRL 10th Computer Networking Conference*, pages 62–66. ARRL, 1991.
- [11] Leonard Kleinrock and John Silvester. Spatial reuse in multihop packet radio networks. *Proceedings of the IEEE*, 75(1):156–167, January 1987.
- [12] H. Vincent Poor. *An introduction to signal detection and estimation*. Springer-Verlag, 2nd edition, 1994.
- [13] Michael B. Pursley. The role of spread spectrum in packet radio networks. *Proceedings of the IEEE*, 75(1):116–134, January 1987.
- [14] S. Ramanathan and Errol L. Lloyd. Scheduling algorithms for multi-hop radio networks. In *SIGCOMM'92 Conference Proceedings*, pages 211–222. ACM, 1992.
- [15] Henry R. Reed and Carl M. Russell. *Ultra High Frequency Propagation*. Boston Technical Publishers, second edition, 1964.
- [16] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, XXVII(3 and 4):379–423 and 623–656, July (part 1) and October (part 2) 1948. Also republished in [20], [21], and [17].
- [17] Claude E. Shannon and Warren Weaver. *The mathematical theory of communication*. University of Illinois Press, 1963.
- [18] Timothy Jason Shepard. *Decentralized Channel Management in Scalable Multihop Spread-Spectrum Packet Radio Networks*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, July 1995. also published as technical report MIT/LCS/TR-670 available on the Internet as <ftp://ftp-pubs.lcs.mit.edu/pub/lcs-pubs/tr.outbox/MIT-LCS-TR-670.ps.gz>.
- [19] Marvin K. Simon, Jim K. Omura, Robert A. Scholtz, and Barry K. Levitt. *Spread Spectrum Communications*. McGraw-Hill, second edition, 1994.
- [20] David Slepian, editor. *Key Papers in The Development of Information Theory*. IEEE Press, 1974.
- [21] N. J. A. Sloane and Aaron D. Wyner, editors. *Claude Elwood Shannon Collected Papers*. IEEE Press, 1993.
- [22] Elvino S. Sousa. Interference modeling in a direct sequence spread-spectrum packet radio network. *IEEE Transactions on Communications*, 38(9):1475–1482, September 1990.
- [23] Elvino S. Sousa and John A. Silvester. Spreading code protocols for distributed spread-spectrum packet radio networks. *IEEE Transactions on Communications*, 36(3):272–281, March 1988.
- [24] Hideaki Takagi and Leonard Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *IEEE Transactions on Communications*, COM-32(3):246–257, March 1984.
- [25] Gregory D. Troxel. Time surveying: Clock synchronization over packet networks. Technical Report MIT/LCS/TR-623, Massachusetts Institute of Technology Laboratory for Computer Science, May 1994.
- [26] Paul F. Tsuchiya. The landmark hierarchy: A new hierarchy for routing in very large networks. In *SIGCOMM Symposium, Stanford, CA*, pages 35–42. ACM, August 1988.
- [27] Sergio Verdú. Recent progress in multiuser detection. In William A. Porter and Subhash C. Kak, editors, *Lecture Notes in Control and Information Sciences*. Springer-Verlag, 1989. Also republished in [2].